

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: Richard Bussiere et al.
Serial No.: 10/713,560
Filed: November 14, 2003
For: DISTRIBUTED INTRUSION RESPONSE SYSTEM
Assignee: Enterasys Networks, Inc.
Examiner: Christopher Brown
Art Unit: 2134 Confirmation No. 8242

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

DECLARATION OF CHRIS CASEIRO PURSUANT TO 37 CFR § 1.131

Dear Sir:

In support of the Applicants' claim of prior invention of the invention described in the referenced application in view of the Sung et al. reference cited in the January 9, 2008, office action, I hereby declare as follows:

1. My name is Chris Caseiro. I am the Applicants' representative in the prosecution of the referenced application.
2. On January 2, 2003, I received from co-inventor Richard Bussiere an email describing the invention that is the subject matter of the referenced application. A copy of the January 2, 2003, email message from Mr. Bussiere and accompanying attachments have previously been submitted as attachment Exhibits A-C to the October 23, 2007, Declaration of Richard Graham.
3. On January 30, 2003, I received from co-inventor David Harrington an email including additional comments related to the invention disclosure of previously described Exhibits B and C of the Graham Declaration. A copy of that January 30, 2003, email from Mr. Harrington to me is attached hereto as Exhibit 1.
4. I have reviewed Exhibit 1 and have compared the information contained in that evidence with currently pending independent Claims 1 and 30 of the present application. I note that all elements from those claims are described in Exhibit 1, including those elements specifically identified in the accompanying second Declaration of Richard Graham. In addition, the following element of independent Claims 1 and 30 are included in Exhibit 1 as follows:

In Claim 1:

“excluding from at least one of the plurality of interconnection devices a policy enforcement module for effecting its own signal transfer policy changes” is described on page 5 of Exhibit 1:

Determine which switch/router along the path has the capabilities needed to prevent the attack effectively. It may be necessary to configure more than one device to stop the attack effectively. It may also (sic) be desirable to NOT stop the attack at the network ingress edge, but rather permit the attack to be redirected to a honeypot, and stopping intrusion into the real network by configuring a device further away from the edge


In Claim 30:

“wherein at least one of the plurality of interconnection devices excludes the policy enforcement module to establish therein the function to change selectively its own signal transfer policies” is described on page 5 of Exhibit 1:

Determine which switch/router along the path has the capabilities needed to prevent the attack effectively. It may be necessary to configure more than one device to stop the attack effectively. It may also (sic) be desirable to NOT stop the attack at the network ingress edge, but rather permit the attack to be redirected to a honeypot, and stopping intrusion into the real network by configuring a device further away from the edge

5. Upon information and belief, the accompanying second Declaration of Richard Graham identifies specifically contemporaneous evidence of the other elements of independent Claims 1 and 30.

6. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 USC 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued.

By: 
Chris Caseiro

Date: May 8, 2008

Certificate of Transmission

I hereby certify that this correspondence is being transmitted to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, using the EFS-Web service of the US Patent Office on May 8, 2008. It is hereby requested that this communication be assigned a receipt date of May 8, 2008.

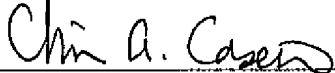

Chris A. Caseiro

Exhibit 1

From: "Caseiro, Chris" <ccaseiro@enterasys.com>
To: <ccaseiro@verrilldana.com>
Date: 1/30/03 2:17PM
Subject: FW: Distributed Intrusion Response System.doc

Re DIRS disclosure 2

-----Original Message-----

From: Harrington, David
Sent: Thursday, January 30, 2003 2:05 PM
To: Caseiro, Chris
Subject: FW: Distributed Intrusion Response System.doc

-----Original Message-----

From: Harrington, David
Sent: Friday, December 20, 2002 1:06 PM
To: Bussiere, Dick
Subject: RE: Distributed Intrusion Response System.doc

Hi Dick,

A few more comments:

The DIRS does this by

Verifying the legitimacy of the originating address might be helped by the NMS, in conjunction with pinging the address directly. The NMS may already know that there is a rogue IP, or that all the devices with the same IP are valid.

Generate an alert to one or more system admins and/or one or more NMS applications.

Can we block ingress based on BOTH the protocol and address, such as SNMP traffic to port UDP 161 originating from IP address a.b.c.d? If the device connected to our ingress port is a shared media device, it may be perfectly acceptable to permit SNMP:161 traffic from other addresses on that same physical ingress port, or to allow other types of traffic from the address a.b.c.d.

components

"; which" is not good English

I'd still like to see L7 mentioned. It is a definite part of our NM plan to have application-traffic recognition, e.g. SAP, and to apply UPN-style filters based on that detection.

topology information should be able to be forwarded to an NM application, which can store the data in a directory. There should not be a need for the devices to communicate directly with a directory. I.e., I don't think we will ever support LDAP in our devices, but we will in our applications. We will support NM protocols, such as SNMP, in

our devices to pass information to NM applications, who can filter, aggregate, and convert data into a format appropriate for sending the data into an LDAP-compatible directory service.

Operational model

Attached device learns upper layer address mappings (L4 to MAC, L7 to MAC, etc.)

You talk about the directory service, and I think you may be thinking of an SFS-style directory. I don't know SFS, but it is my impression that an SFS-style directory is not currently in our plans. This would require further discussion.

It might be useful to provide more than just the IP address to directory service, if more is known. We may have lots of information via Node/Alias and RMON about the addresses and the traffic patterns associated with those addresses. This information might help eliminate false positives, or help to identify a pattern of attack behavior through the relevant access device.

I think it important to include in here a discussion of the fact that the Policy Manager assesses the request and determines whether to apply the requested change. It should not be assumed by the IRS that the requested change has been made.

I think this document needs to briefly discuss the expectations of a NM application (Netsight or 3rd party) to aggregate topology information, and the expectations that a policy manager can receive requests to modify policies, and the expectations that a policy manager will need to somehow adjudicate whether the policy change request should be granted based on its knowledge of the network and administrator-defined policies about such requests. These may be phase 2 or 3 deliverables, but the expectations need to be identified since the policy management will be an important part of the DIRS system.

dbh

-----Original Message-----

From: Bussiere, Dick

Sent: Thursday, December 19, 2002 3:12 PM

To: Harrington, David

Subject: RE: Distributed Intrusion Response System.doc

Hi,

Attached is a slightly revised version, based on your comments. Thank you for taking the time to review it. Please see below for more questions/comments, in [...]

-----Original Message-----

From: Harrington, David

Sent: Tuesday, December 17, 2002 3:50 PM

To: Bussiere, Dick

Subject: RE: Distributed Intrusion Response System.doc

Hi Dick,

I think this is very good. There are a few places where I might use a slightly different approach. Comments inline.

Do you want to work together to get this fleshed out?

dbh

-----Original Message-----

From: Bussiere, Dick

Sent: Tuesday, December 17, 2002 1:21 PM

To: Harrington, David

Subject: Distributed Intrusion Response System.doc

Please read this and tell me what you think.

Thank you,

-d

Distributed Intrusion Response System

The Distributed Intrusion Response System (DIRS) is a plurality of components which is intended to cause a dynamic response to a detected intrusion. The purpose of the DIRS is to:

I'm not crazy about DIRS as the name since it sound more like directory than IDS or IRS. Maybe D-IRS would be better, or

Policy-based Intrusion Response (PBIR).

[I have left it as DIRS for the time being. The marketing people can give it a clever name.]

* Stop a detected attack in a very granular fashion by leveraging policy enforcement controls available at the edge of the IT infrastructure

I would have said fine-grained. Very granular can also be interpreted as coarse, as in a grainy photograph or granular sugar vs. powdered.

[Changed as suggested]

* Prevent detected attacks from spreading, through an early reaction mechanism

I added a comma; it sounded like the early reaction mechanism was the tool that permitted the attack to spread.

[Changed as suggested]

* Only disable protocols which are active participants in the attack, while leaving other protocols enabled

I think one of the important things to shoot for might be look-ahead; given the protocol signatures already detected, which protocol signatures do we need to try to block to prevent the rest of the attack signature.

[Wording changed as suggested]

* Dynamically disable protocols which are abused, leaving other productive protocols enabled

* Minimize effects of false positives

The DIRS does this by:

- * Detecting a probable attack through 'traditional' IDS methods, either host based or network based
- * Identifying the source of the attack
- * Determining where in the network topology the attack was originating from
- * Verifying the legitimacy of the attack origin (prevent IP address spoofing attacks from causing denial of service to legitimate end users)
- * If available, determine WHO is generating the attack

Determine the path from the originating point to the target system

Determine which switch/router along the path has the capabilities needed to prevent the attack effectively. It may be necessary to configure more than one device to stop the attack effectively. It may also be desirable to NOT stop the attack at the network ingress edge, but rather to permit the attack to be redirected to a honeypot, and stopping intrusion into the real network by configuring a device further away from the edge.

[I think that this topology stuff would be good as a phase II; it sounds like it may be quite complicated to implement. For now I would be happy just to leverage the capabilities of the edge.]

- * Block the attack at the point of entry; either the specific protocol or the source MAC/IP address

or the protocol and the associated protocol-specific address

- * Generate an alert to the system administrator

The components of the DIRS are:

- * A centralized or distributed Intrusion detection system or systems; which monitor the network or networks for malicious or potentially malicious activities

This might be especially effective if hardware-assisted IDS functionality could be co-located within the edge switch/router.

[There are plans afoot to do this, we intend to architect a multi-purpose IDS bolt-on.]

* Distributed network security policy enforcement devices which are capable of enforcing L2, L3, and L4 policies in any combination

and L7?

[Where practical.]

* The same network security policy devices (discussed above) having the ability to generate L2 location information, L2-L3 mapping information, and having the ability to forward this information to a directory service (a distributed directory)

* A directory service which can be used to locate the location within the topology in which a given physical device is located

I would reword this. "locate the location where a device is located" seems to be a poorly constructed sentence.

[Done]

Maybe "A directory service to provide the location of a physical device within the network topology."

[Done]

I think it is important to understand the location within the physical topology, but also within the logical topologies of the network. Understanding the physical location can be important for trying to catch the intruder, or to correlate multiple different types of attacks coming through the same ingress point to the network, and for assessing whether the detected behavior is coming from a trusted or untrusted zone, e.g. wireless from <guest> in the parking lot or wired from <authenticated operator> in the NOC."

Understanding the device position within the logical

topology is important to prevent applying policies in bad ways; shutting down spanning tree BPDUs at a key switching point, or BGP updates at a border gateway device would be a horrible response. It is important to understand what role the device plays in the functioning of the network before blindly applying automated policy rules.

[If we only allow access ports to be modified then is this a problem?]

This points out one piece that is missing from your system - the policy manager that allows the administrators to set policies that vary according to the physical location and logical role of the device.

I am also concerned that the IDS is sending out control information directly to the devices. It is critically important to maintain network stability and integrity. We sell a policy management system; we don't want the two systems generating conflicting policy directions; I recommend the IDS request the policy-management system to modify the policy applied to the physical port or logical channel. This way the policy-management system knows just what policies are being applied to the network by the policy enforcement points.

[Check out the drawing and wording. I agree with you on this and I have addressed it, I think.]

One of the very important things being discussed in the O&M Area of the IETF is the fact that operators do not want devices tweaked in unpredictable ways by policy-management systems. I think it will be important to make it possible to define a small set of tested policies that can be applied to a user, and when certain conditions are met, they are forced to reauthenticate, and their authorization may be changed to a more limiting access template. This is very consistent with the UPN approach.

These components are shown in the illustration below:

IETF standardized terminology for policy-based

management includes Policy Enforcement Points (PEPs) which may or may not be devices, and Policy Decision Points (PDPs) which determine which policies should be distributed to the PEPs for enforcement.

I would draw the picture with the policy-based network management system (the PDP) using the same or a supplementary directory service to determine network configuration policies and distributing them to the policy enforcement points (PEPs).

The IDS should send control requests to the PDP, which could modify the policies for the offender and distribute the modified policies to the PEPs. The policies should reflect the administrators' preferences for allowed responses based on the user roles and the device roles within the organization.

I would also recommend that we consider building in case-based reasoning into the system, something contained in an Aprisma patent we are licensed to use. This is a mechanism that asks an administrator to craft a new response to be learned by the system, or to select which known response they want applied to the current situation, and asks the administrator if they want the same response used automatically in the future when a similar situation is detected. (emphasis provided by Outlook not under my control).

This allows the administrator to try the response without making it an automated response, and allows the administrator to change the preferences to make the response automatic once they are satisfied it will not impact network stability. A CBR solution could also allow an existing response to be modified to meet additional conditions, or to not take effect under other conditions, as the system is taught how to respond on a case-by-case basis.

Operational Model briefly described:

- * Policy enforcement device "learns"
- * Attached device L2 physical (MAC)
address
- * Attached device L3 logical to physical
(i.e. MAC to IP) address binding
- plus other protocol-specific address/MAC bindings

- * Policy Enforcement Device reports learned topology information to directory service; included in this information is information relating to specific identification of Policy Enforcement Device (i.e. IP address of switch or router)
- * Intrusion Detection Device monitors network for malicious activities
- * On detection of malicious activity, if Intrusion Detection Device policy for the activity indicates that Intrusion Response is required, Intrusion Detection Device contacts Directory Service with source IP address of captured packet
- * Directory Service replies with L2 (MAC address) and information sufficient to identify Policy Enforcement Device to which the attacker, as identified by L3 IP address, is attached
- * IDS will determine if IP address is legitimate by testing for presence of it

Is this something that IDS already is expected to do? or is this something an IRS, in conjunction with the NMS, should be responsible for?

[People spoof addresses, so we need to figure out if an address is legit before we terminate it.]

- * IDS modifies policy in Policy Enforcement Device, by sending control information to the device (i.e. SNMPv3) such that:

Again, to the PDP, who assesses whether to honor the request, and responds to the IDS with information about whether the request was granted. The IDS/IRS at this point could request different mechanisms to try to stop the attacks. Given a detected signature that had an SNMP attack, and an anticipated HTTP attack, the PDP may refuse a request to deny all SNMP traffic, but might grant a request to block all HTTP traffic.

If the PDP changes the policy, it will need to keep track of the original policy and the new policy and alert the network administrator and probably the security administrator as well.

The PDP has knowledge of what protocols are needed to run the network to meet business objectives. The IRS knows what the signature of the attack looks like, and knows how to prevent the anticipated attack, but doesn't know about what is needed to keep the business operating properly. The IRS should work through the PDP (which actually might work through the NMS to deploy policies).

- * Further network access by the device is blocked

completely (i.e. MAC address filter is installed

- * Access by the IP address (only) is blocked (L3 filter installed)

- * Access or use of the offending protocol is blocked (L2/L4 or L3/L4 protocol is blocked)

Novel Concepts:

The concept of an intrusion detection system generating an active response to a detected event is not new. For example, Dragon does this with the "shun" and "snipe" features. Unfortunately, the current implementation is static in nature in that it cannot take into account network topology. The device which will perform the shunning, for example, must be known in advance and "hard coded".

The proposed approach is novel in that no prior knowledge of the actual network element which will be used to block the suspected traffic is required. This means that the number of policy enforcement points is infinite; any network element can be thought of as a device which is capable of enforcing security policy. This in effect "distributes" the intrusion response mechanism - creating in effect a "web". On detection of an attack, the network element which services the attacking device is learned from the directory service. The network element is then reconfigured to modify the authorization for the offending client machine.

Possibly patentable ideas include:

- * Apparatus for accumulating topology information in an intrusion detection system

The IRS might still gather topology info from the NMS to assess how best to stop the attack.

- * Apparatus for determining connection point of end station within an intrusion detection system

- * Mechanism for automating intrusion response mechanism

And we can supplement these with:

apparatus for sharing topology information with IDS

apparatus for identifying connection point to an IDS

apparatus for responding to an IRS request, taking into consideration network stability

apparatus for learning administrator preferences for intruder responses using CBR

Implementation Ideas:

The current concept is to leverage the existing NetSight "Compass" application as the directory service.

I would leverage the information made available by the Compass application, not necessarily the Compass application itself. If Compass can store the information in an SQL database, the IRS can query the database as needed. Compass accumulates the information from Enterasys switches. Compass would be bundled with Dragon as an option which would enable this functionality. Some development work would need to be done to give Compass an API which Dragon could leverage and use to query the directory.

Dragon should have a flexible API which would allow it to "plug-in" to applications other than Compass which are capable of collecting directory and network topology information.

I recommend a data model in the IRS that can be mapped to the data model of a SQL database (like synching a Palm pilot to an Outlook or Netscape address book).

It may also be quite interesting to develop a Compass plug-in which would allow it to be extended to support third-party switches which are capable of collecting L2, L3 and L4 information and enforcing L2, L3 and L4 policies. Yes, and to use RMON to collect information about devices.

It would also be good if the requests made by the IRS were compatible with standards-based approaches, such as SNMPCONF, which could be applied to any switch that supported SNMPCONF and the desired functionality.

dbh